APPENDIX I

TCP/IP Illustrated, Volume 1

The Protocols

W. Richard Stevens

ADDISON-WESLEY

Boston • San Francisco • New York • Toronto • Montreal
London • Munich • Paris • Madrid
Capetown • Sydney • Tokyo • Singapore • Mexico City

We have removed all the window advertisements from this figure, except for segment 72, which we discuss below. `slip` always advertised a window of 4096, and `vangogh` advertised a window of 8192. The segments are numbered in this figure as a continuation of Figure 21.2, where the first data segment across the connection was numbered 1. As in Figure 21.2, the segments are numbered according to their send or receive order on the host `slip`, where `tcpdump` was being run. We have also removed a few segments that have no relevance to the discussion (44, 47, and 49, all ACKs from `vangogh`).

It appears that segment 45 got lost or arrived damaged—we can't tell from this output. What we see on the host `slip` is the acknowledgment for everything up through but not including byte 6657 (segment 58), followed by eight more ACKs of this same sequence number. It is the reception of segment 62, the third of the duplicate ACKs, that forces the retransmission of the data starting at sequence number 6657 (segment 63). Indeed, Berkeley-derived implementations count the number of duplicate ACKs received, and when the third one is received, assume that a segment has been lost and retransmit only one segment, starting with that sequence number. This is Jacobson's *fast retransmit* algorithm, which is followed by his *fast recovery* algorithm. We discuss both algorithms in Section 21.7.

Notice that after the retransmission (segment 63), the sender continues normal data transmission (segments 67, 69, and 71). TCP does not wait for the other end to acknowledge the retransmission.

Let's examine what happens at the receiver. When normal data is received in sequence (segment 43), the receiving TCP passes the 256 bytes of data to the user process. But the next segment received (segment 46) is out of order: the starting sequence number of the data (6913) is not the next expected sequence number (6657). TCP saves the 256 bytes of data and responds with an ACK of the highest sequence number successfully received, plus one (6657). The next seven segments received by `vangogh` (48, 50, 52, 54, 55, 57, and 59) are also out of order. The data is saved by the receiving TCP, and duplicate ACKs are generated.

Currently there is no way for TCP to tell the other end that a segment is missing. Also, TCP cannot acknowledge out-of-order data. All `vangogh` can do at this point is continue sending the ACKs of 6657.

When the missing data arrives (segment 63), the receiving TCP now has data bytes 6657–8960 in its buffer, and passes these 2304 bytes to the user process. All 2304 bytes are acknowledged in segment 72. Also notice that this ACK advertises a window of 5888 (8192 – 2304), since the user process hasn't had a chance to read the 2304 bytes that are ready for it.

If we look in detail at the `tcpdump` output for the dips around times 14 and 21 in Figure 21.6, we see that they too were caused by the receipt of three duplicate ACKs, indicating that a packet had been lost. In each of these cases only a single packet was retransmitted.

We'll continue this example in Section 21.8, after describing more about the congestion avoidance algorithms.

2. A set of common structures and an identification scheme used to ref variables in the MIB. This is called the *Structure of Management Informa* and is specified in RFC 1155 [Rose and McCloghrie 1990]. For exampl specifies that a `Counter` is a nonnegative integer that counts from 0 4,294,967,295 and then wraps around to 0.
3. The protocol between the manager and the element, called the *Simple Management Protocol* (SNMP). RFC 1157 [Case et al. 1990] specifies the This details the format of the packets exchanged. Although a wide v transport protocols could be used, UDP is normally used with SNMP.

These RFCs define what is now called SNMPv1, or just SNMP, which is the topi chapter. During 1993 additional RFCs were published specifying SNMP Ve (SNMPv2), which we describe in Section 25.12.

Our approach to SNMP in this chapter is to describe the protocol between th ager and the agent first, and then look at the data types for the variables maintai the agent. We describe the database of information maintained by the agent (the looking at the groups that we've described in this text: IP, UDP, TCP, and so o show examples at each point along the way, tying network management back protocol concepts from earlier chapters.

25.2 Protocol

SNMP defines only five types of messages that are exchanged between the manag agent.

1. Fetch the value of one or more variables: the `get-request` operator.
2. Fetch the next variable after one or more specified variables: `get-next-request` operator. (We describe what we mean by "next" lat this chapter.)
3. Set the value of one or more variables: the `set-request` operator.
4. Return the value of one or more variables: the `get-response` operator. Th the message returned by the agent to the manager in response to `get-request`, `get-next-request`, and `set-request` operators.
5. Notify the manager when something happens on the agent: the `trap` operato

The first three messages are sent from the manager to the agent, and the last two from the agent to the manager. (We'll refer to the first three as the `get`, `get-next`, a `set` operators.) Figure 25.1 summarizes these five operators.

Since four of the five SNMP messages are simple request-reply protocols (the ma ager sends a request, the agent sends back a reply) SNMP uses UDP. This means that request from the manager may not arrive at the agent, and the agent's reply may n make it back to the manager. The manager probably wants to implement a timeout an retransmission.

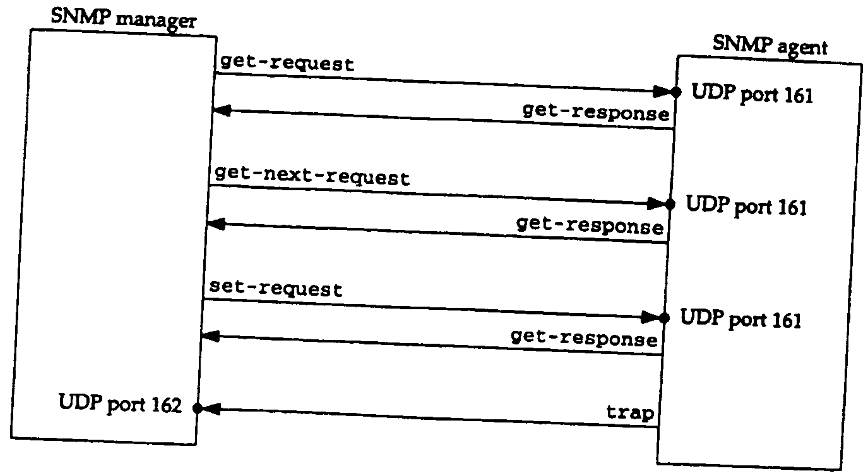

Figure 25.1 Summary of the five SNMP operators.

The manager sends its three requests to UDP port 161. The agent sends traps to UDP port 162. By using two different port numbers, a single system can easily run both a manager and an agent. (See Exercise 25.1.)

Figure 25.2 shows the format of the five SNMP messages, encapsulated in a UDP datagram.

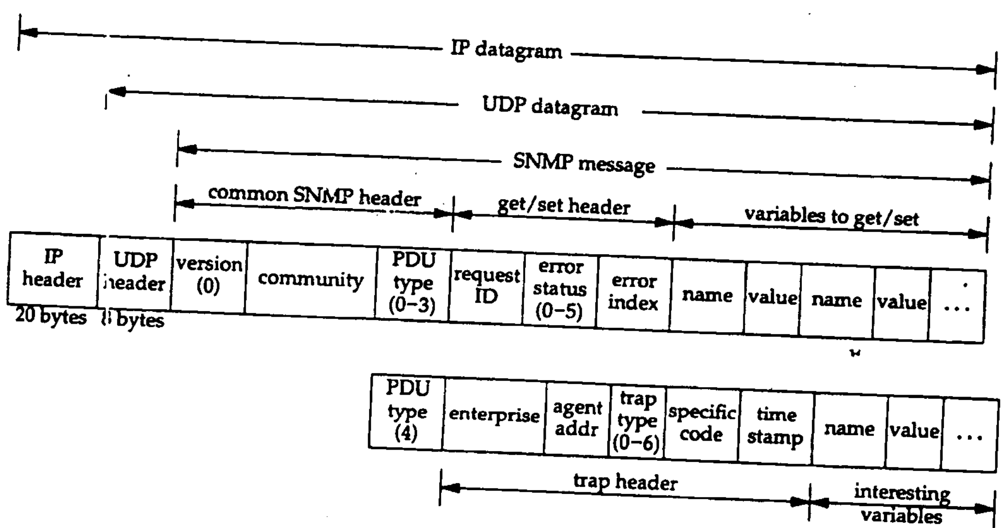

Figure 25.2 Format of the five SNMP messages.

In this figure we specify the size in bytes of the IP and UDP headers only. This is because the encoding used for the SNMP message—called ASN.1 and BER, which we

```
sun % snmpi -a netb -c secret get ipRouteNextHop.140.252.3.0
ipRouteNextHop.140.252.3.0=140.252.1.6
```

This routing table entry tells netb to send the packets to swnrt, which is what we see happen.

Why does gemini send the packets directly back through netb? Because on gemini the destination address of the return packets is 140.252.1.29, and that network (140.252.1) is on a directly connected interface.

What we're seeing in this example is a policy routing decision. The default route to network 140.252.3 is through the router swnrt because gemini is intended to be a multihomed host, not a router. This is an example of a multihomed host that does not want to be a router.

Traps

All the examples we've looked at so far in this chapter have been from the manager to the agent. As shown in Figure 25.1, it's also possible for the agent to send a *trap* to the manager, to indicate that something has happened on the agent that the manager might want to know about. Traps are sent to UDP port 162 on the manager.

In Figure 25.2 we showed the format of the trap PDU. We'll go through all the fields in this message when we look at some tcpdump output below.

Six specific traps are defined, with a seventh one allowing a vendor to implement an enterprise-specific trap. Figure 25.30 describes the values for the *trap type* in the trap message (Figure 25.2).

trap type	Name	Description
0	coldStart	Agent is initializing itself.
1	warmStart	Agent is reinitializing itself.
2	linkDown	An interface has changed from the up to the down state (Figure 25.18). The first variable in the message identifies the interface.
3	linkUp	An interface has changed from the down to the up state (Figure 25.18). The first variable in the message identifies the interface.
4	authenticationFailure	A message was received from an SNMP manager with an invalid community.
5	egpNeighborLoss	An EGP peer has changed to the down state. The first variable in the messages contains the IP address of the peer.
6	enterpriseSpecific	Look in the *specific code* field for information on the trap.

Figure 25.30 Trap types.

We can see some traps using tcpdump. We'll start the SNMP agent on the system and see it generate a coldStart trap. (We tell the agent to send traps to the host .) Although we're not running a manager on bsdi to handle the traps, we can run

tcpdump and see what packets get generated. Recall from Figure 25.1 that a tr
from the agent to the manager, but there is no acknowledgment sent by the ma
we don't need a manager to handle the traps.) We then send a request using th
program, but with an invalid community name. This should gen
authenticationFailure trap. Figure 25.31 shows the output.

```
1  0.0            sun.snmp > bsdi.snmp-trap: C=traps Trap(28)
                  E:unix.1.2.5 [140.252.13.33] coldStart 20
2  18.86 (18.86)  sun.snmp > bsdi.snmp-trap: C=traps Trap(29)
                  E:unix.1.2.5 [140.252.13.33] authenticationFailu
```

Figure 25.31 tcpdump output of traps generated by SNMP agent.

First we notice that both UDP datagrams are from the SNMP agent (po
printed as the name snmp) with a destination port of 162 (printed as the
snmp-trap).

The notation C=traps is the community name of the trap message. This is
figuration option with the ISODE SNMP agent being used.

The next notation, Trap(28) in line 1 and Trap(29) in line 2 is the *PDU typ*
length.

The next field of output for both lines is E:unix.1.2.5. This is the *enterpri*
agent's sysObjectID. It falls under the 1.3.6.1.4.1 node of the tree in Figure
(iso.org.dod.internet.private.enterprises), so this agent's object iden
is 1.3.6.1.4.1.4.1.2.5. Its abbreviated name is unix.agents.fourBSD-isode.5.
final number (5) is the version number of this release of the ISODE agent. This en
prise value identifies the agent software generating the trap.

The next field output by tcpdump is the IP address of the agent (140.252.13.33).

The *trap type* is printed as coldStart on line 1, and authenticationFailu
on line 2. These correspond to *trap type* values of 0 and 4, respectively (Figure 25.
Since these are not enterprise-specific traps, the *specific code* must be 0, and is
printed.

Next comes the *timestamp* field, printed as 20 and 1907. This is a TimeTicks valu
representing the number of hundredths of a second since the agent initialized. In t
case of the cold start trap, the trap was generated 200 ms after the agent was initializ
The tcpdump output indicates that the second trap occurred 18.86 seconds after the fir
one, which corresponds to the printed value of 1907 hundredths of a second, minus 20
ms.

Figure 25.2 indicates that a trap message can contain interesting variables that th
agents wants to send to the manager, but there aren't any in our examples.

25.11 ASN.1 and BER

The formal specification of SNMP uses *Abstract Syntax Notation 1* (ASN.1) and the actual
encoding of the bits in the SNMP messages (Figure 25.2) uses the corresponding *Basic*
Encoding Rules (BER). Unlike most texts that describe SNMP, we have purposely left a

discussion
the reader
tion we o
ASN.1 an
ASN.1
says nothi
SNMP me
data type

IpAdd
[

Similarly,

udpNc

The defi
Give
stream o
such as
byte say
the bina
Fort
SNMP.
ment.

25.12 SNMP

During
RFC 14
Two bo
implem
tations
In

1.

2.

APPENDIX J

UNITED STATES DEPARTMENT OF COMMERCE
Patent and Trademark Office
Address: COMMISSIONER OF PATENTS AND TRADEMARKS
Washington, D.C. 20231

APPLICATION NUMBER	FILING DATE	FIRST NAMED APPLICANT	ATTORNEY DOCKET NO.
08/644,072	05/09/96	MOR[illegible]	

PENNIE & EDMONDS
1155 AVENUE OF THE AMERICAS
NEW YORK NY 10036-2711

EXAMINER

ART UNIT	PAPER NUMBER

DATE MAILED:

This is a communication from the examiner in charge of your application.
COMMISSIONER OF PATENTS AND TRADEMARKS

OFFICE ACTION SUMMARY

☐ Responsive to communication(s) filed on ______

☐ This action is FINAL.

☐ Since this application is in condition for allowance except for formal matters, **prosecution as to the merits is closed** in accordance with the practice under *Ex parte Quayle*, 1935 D.C. 11; 453 O.G. 213.

A shortened statutory period for response to this action is set to expire 3 month(s), or thirty days, whichever is longer, from the mailing date of this communication. Failure to respond within the period for response will cause the application to become abandoned. (35 U.S.C. § 133). Extensions of time may be obtained under the provisions of 37 CFR 1.136(a).

Disposition of Claims

☑ Claim(s) 1–37 is/are pending in the application.

Of the above, claim(s) ______ is/are withdrawn from consideration.

☐ Claim(s) ______ is/are allowed.

☑ Claim(s) 1-11, 13, 14, 16-22, 24, 25, 26-36 is/are rejected.

☑ Claim(s) 12, 15, 23, 37 is/are objected to.

☐ Claims ______ are subject to restriction or election requirement.

Application Papers

☑ See the attached Notice of Draftsperson's Patent Drawing Review, PTO-948.

☐ The drawing(s) filed on ______ is/are objected to by the Examiner.

☐ The proposed drawing correction, filed on ______ is ☐ approved ☐ disapproved.

☐ The specification is objected to by the Examiner.

☐ The oath or declaration is objected to by the Examiner.

Priority under 35 U.S.C. § 119

☐ Acknowledgement is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d).

☐ All ☐ Some* ☐ None of the CERTIFIED copies of the priority documents have been

☐ received.

☐ received in Application No. (Series Code/Serial Number) ______

☐ received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

*Certified copies not received: ______

☐ Acknowledgement is made of a claim for domestic priority under 35 U.S.C. § 119(e).

Attachment(s)

☑ Notice of Reference Cited, PTO-892

☐ Information Disclosure Statement(s), PTO-1449, Paper No(s). ______

☐ Interview Summary, PTO-413

☑ Notice of Draftsperson's Patent Drawing Review, PTO-948

☐ Notice of Informal Patent Application, PTO-152

– SEE OFFICE ACTION ON THE FOLLOWING PAGES –

PTOL-326 (Rev. 10/95)

08/644072

Claim Rejections - 35 USC § 102

1. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless --
>
> (b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

2. Claims 1-11, 16-22, 26-33 are rejected under 35 U.S.C. 102(b) as being clearly anticipated by U.S. Patent No. 5,604,542 to Dedrick.

As per claims 1, 16, 26 and 27, Dedrick discloses converting a plurality of streams of audio and/or visual information into a plurality of streams of digital packets (col. 2, lines 38-40), routing the stream to one or more users (col. 3, lines 16-18), controlling the routing of the packets in response to selection signals from the users (col. 3, lines 58-60) and monitoring the reception of packets by the users and accumulating records (col. 4, lines 45-48).

As per claims 2, 8, 17, 20, 28 and 32, Dedrick discloses at least some advertising in at least one stream of packets (col. 2, lines 12-14).

As per claims 3, 9, 18, 29 and 33, Dedrick discloses varying the content of the advertising information with the identity of

the user to whom the advertising information is provided (col. 3, line 65 to col. 4, line 3).

As per claims 4 and 7, Dedrick discloses the information being inserted the stream before the stream is converted into a stream of packets (col. 2, lines 12-14).

As per claims 5 and 30, Dedrick discloses generating an audio output and/or a visual display from the stream of packets (col. 3, lines 33-37).

As per claims 6, 19 and 31, Dedrick discloses storing a first stream at a first time and inserting the first stream in a second stream received by the user (col. 3, line 67 to col. 4, line 3).

As per claims 10, 11, 21 and 22, Dedrick discloses the records indicate how many users received specific advertising information (col. 4, lines 45-47), this information being music selections (col. 2, 18-20).

Claim Rejections - 35 USC § 103

3. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

> (a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

4. Claims 13, 14, 24, 25, 34, 35 and 36 are rejected under 35 U.S.C. 103(a) as being unpatentable over Dedrick.

As per claims 13, 14, 24, 25, 34, 35 and 36, Dedrick does not disclose the limitations of these claims. It is known to compress information at the source and decompress it at the destination, selecting the compression algorithm in accordance with the content of the information. It would have been obvious to a person having an ordinary level of skill in the art at the time the invention was made to have modified the system disclosed by Dedrick to include these features since they allow for efficient transmission of data at a faster rate using less memory storage space and selection of the compression algorithm ensures a reliable reception with little loss of quality and minimum degradation.

Allowable Subject Matter

5. Claims 12, 15, 23, and 37 objected to as being dependent upon a rejected base claim, but would be allowable if rewritten in independent form including all of the limitations of the base claim and any intervening claims.

Conclusion

6. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

U.S. Patent No. 5,493,514 to Keith et al. discloses Process, Apparatus and System for Encoding and Decoding Video Signals.

7. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Mr. Thomas Peeso whose telephone number is (703) 305-9784. The examiner can normally be reached on Monday -Thursday from 7am to 5pm. The examiner can also be reached on alternate Fridays.

8. If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Mr. Emanuel Voeltz, can be reached on (703) 305-9714. The fax phone number for this Group is (703) 305-95[64,65].

9. Any inquiry of a general nature or relating to the status of this application should be directed to the Group receptionist whose telephone number is (703) 305-3800.

Thomas Peeso
Patent Examiner
Art Unit 2414
15 Jul 97

FORM PTO-892 (REV. 2-92)

U.S. DEPARTMENT OF COMMERCE
PATENT AND TRADEMARK OFFICE

SERIAL NO. 08/644,672 | GROUP ART UNIT 2414 | ATTACHMENT TO PAPER NUMBER 7

NOTICE OF REFERENCES CITED

APPLICANT(S): MONTIERO ET AL.

U.S. PATENT DOCUMENTS

*		DOCUMENT NO.	DATE	NAME	CLASS	SUB-CLASS	FILING DATE IF APPROPRIATE
	A	5604542	18Feb97	Dedrick	348	552	8Feb95
	B	5493514	20Feb96	Keith et al.	364	514R	
	C						
	D						
	E						
	F						
	G						
	H						
	I						
	J						
	K						

FOREIGN PATENT DOCUMENTS

*		DOCUMENT NO.	DATE	COUNTRY	NAME	CLASS	SUB-CLASS	PERTINENT SHTS. DWG.	PERTINENT PP. SPEC.
	L								
	M								
	N								
	O								
	P								
	Q								

OTHER REFERENCES (Including Author, Title, Date, Pertinent Pages, Etc.)

R	
S	
T	
U	

EXAMINER: T. PEESO | DATE: 14 JUL 97

* A copy of this reference is not being furnished with this office action.
(See Manual of Patent Examining Procedure, section 707.05 (a).)